RIACS

Research Institute for Advanced Computer Science
NASA Ames Research Center

NASA-CR-200830

7N-99-CR

RIACS / USRA

Annual Report

1995

RIACS Annual Report 1995

Contract NAS 2-13721

Submitted to:

Contracting Officer: Catharine P. Levine
Code IC, MS 241-1

CO, Code ASR, MS 241-1

External Relations Office
Technology Utilization, Code DXR, MS 223-3

Ames Research Patent Counsel
Code DP, MS 200-11A

Submitted by:
Research Institute for Advanced Computer Science (RIACS)

An Institute of:
Universities Space Research Association (USRA)

RIACS Principal Investigator
Joseph Oliger

NASA Technical Monitor
Henry Lum

COTR
Sonie Lau

Joseph Oliger

Joseph Oliger, Director

Table of Contents

I. INTRODUCTION

Joseph Oliger, Director

The Research Institute for Advanced Computer Science (RIACS) was established by the Universities Space Research Association (USRA) at the NASA Ames Research Center (ARC) on June 6, 1983. RIACS is privately operated by USRA, a consortium of universities with research programs in the aerospace sciences, under contract with NASA.

The primary mission of RIACS is to provide research and expertise in computer science and scientific computing to support the scientific missions of NASA ARC. The research carried out at RIACS must change its emphasis from year to year in response to NASA ARC's changing needs and technological opportunities. A flexible scientific staff is provided through a university faculty visitor program, a post doctoral program, and a student visitor program. Not only does this provide appropriate expertise but it also introduces scientists outside of NASA to NASA problems. A small group of core RIACS staff provides continuity and interacts with an ARC technical monitor and scientific advisory group to determine the RIACS mission. RIACS activities are reviewed and monitored by a USRA advisory council and ARC technical monitor.

Research at RIACS is currently being done in the following areas:

- Parallel Computing
- Advanced Methods for Scientific Computing
- High Performance Networks
- Learning Systems

Parallel compiler techniques, adaptive numerical methods for flows in complicated geometries and optimization have been identified as important problems to investigate for ARC's involvement in the Computational Grand Challenges of the next decade.

During the past year Professor Antony Jameson of Princeton University and Professor Wai-Pai Tang of the University of Waterloo have been visiting RIACS.

In 1995 RIACS had 6 staff scientists, ten visiting scientists, four post-doctoral scientists, ten consultants, three research associates and six summer student visitors.

RIACS technical reports are usually preprints of manuscripts that have been submitted to research journals or conference proceedings. A list of these reports for the period January 1, 1995 through December 31, 1995 is in the Reports and Abstracts section of this report.

II. RESEARCH IN PROGRESS

A. PARALLEL COMPUTING

HIGH PERFORMANCE FORTRAN

Robert S. Schreiber

The HPF standard, developed by the HPF Forum in 1992--1993, is now regarded as the best method for data parallel programming. Good implementations on a number of parallel machines are just now appearing.

In 1995 an effort to develop a second generation HPF was begun. Schreiber is again a leading participant, and is the head of the subgroup of the HPF Forum developing language extensions for task parallelism and computation mapping. Among the features being developed by this subgroup are multiple HPF threads, computation mapping through an ON_HOME directive, reduction operations in independent loops, and asynchronous I/O. Other subgroups are extending the class of data mappings allowed in HPF, and in interoperability with C. These new features will add substantially to the class of parallel computations that HPF can express.

COMPILING ARRAY PARALLELISM FOR DISTRIBUTED-MEMORY MACHINES

Robert S. Schreiber, Thomas J. Sheffler, Leonid Oliker (U. of Colorado at Boulder)
John Gilbert (Xerox PARC), Ender Bilir (Princeton University; Summer Intern)

The goal of this project is to develop compiler algorithms to automate the task of data lay out when compiling array-parallel languages such as Fortran 90 on distributed-memory parallel computers. Data layout is accomplished by alignment of arrays to a template and distribution of the template to the processors.

We previously developed a suite of both top-down and bottom-up heuristics for identifying phases in a program that should be performed without data redistribution. This year we developed the bottom-up approach further. We find subgraphs of the program data flow graph that can be shown to be redistribution free in an optimal distribution solution. (The dataflow graph we use is the alignment-distribution graph we have developed to model data parallel programs and their communication costs). We can thus contract these subgraphs to single nodes without loss of quality in the final distribution that results. Our experiments this year verified that the contraction approach is extremely useful in distribution optimization, as it is usually able to reduce problem size by over 90%; in some cases it solves the problem outright by verifying that a static distribution is optimal for the whole computation. We have developed a robust software implementation of this optimization.

Unstructured communication is an expensive data parallel operation. It has previously been shown that an inspector/executor or runtime compilation paradigm is appropriate: one finds reused unstructured communications, and does as much of the housekeeping work as possible in a preprocessing step, saving the generated information for later use in actually performing the communication. We have noted that existing libraries (PARTI, e.g.) for this do not offer as much of a performance advantage as one might wish; moreover, these tools are not usable from HPF. With Ender Bilir's work in the summer we implemented an HPF callable inspector executor for the unstructured gather operation X = Y(V) (V is a vector-valued subscript) in HPF. The performance we achieved was very good -- roughly double that of PARTI, and of the PGI HPF compiler which also uses an inspector/executor approach.

A PARALLEL C++ TEMPLATE LIBRARY
Thomas J. Sheffler

The goal of this project is to develop a portable template library of parallel algorithms. Templates in C++ are subroutines that are parameterized by type. Using templates, it is possible to write generic algorithms that work on any user-defined type of object. Template algorithms can improve programmer productivity by providing standard algorithms that can be used in many contexts. With such a library, a programmer is freed from re-inventing standard algorithms.

This project resulted in the Amelia Vector Template Library (AVTL). It provides a single distributed data type, the vector, and a large number of standard parallel algorithms on vectors. These include such operations as sorting their contents, or computing prefix sums. Because the vector class and all algorithms are templates, vectors of any user-defined type may be created and manipulated.

The AVTL was well-received in the parallel C++ research community. A chapter describing its use will appear in a book to be entitled "Parallel Programming Using C++," to be published by MIT Press sometime in 1996.

DYNAMIC LOAD BALANCING FOR UNSTRUCTURED ADAPTIVE GRIDS
Rupak Biswas, Andrew Sohn (New Jersey Institute of Technology),
Horst D. Simon (Silicon Graphics), and Leonid Oliker (Univ of Colorado)

The computational requirements for an adaptive solution of unsteady problems in fluid dynamics change as the simulation progresses. This causes workload imbalance among processors on a parallel machine which, in turn, requires significant data movement at runtime. If computational fluid dynamics is to be successful on distributed-memory machines, balancing the runtime load is indispensable. A new dynamic load-balancing framework, called JOVE, has been developed that balances the workload across all processors with a global view. Whenever the computational mesh is adapted, JOVE is activated to eliminate the load imbalance. It first determines whether the new mesh should be repartitioned. If it decides to accept the new partitioning, the application is interrupted to redistribute the data as dictated by the mapping of partitions to processors.

JOVE introduces three novel features. First, a dual graph representation of the computational mesh is used to keep the complexity and connectivity constant during the course of an adaptive computation. Second, a new inertial spectral mesh partitioning method is used that performs both faster and better than Recursive Spectral Bisection. Finally, accurate metrics for the computational gain, the communication cost, and the redistribution overhead are developed to identify and balance the processor loads between successive adapted grids.

JOVE has been implemented on a SP2 in MPI for portability. Preliminary results on model meshes indicates that this approach to dynamic load balancing can be effective for large-scale scientific computations on distributed-memory multiprocessors.

B. ADVANCED METHODS FOR SCIENTIFIC COMPUTING

DYNAMIC MESH ADAPTATION OF UNSTRUCTURED HEXAHEDRAL GRIDS

Rupak Biswas

Anisotropic mesh adaptation is a powerful tool for computing steady and unsteady three-dimensional problems that require local grid modifications in order to efficiently resolve solution features. However, repeated anisotropic subdivision can significantly deteriorate the quality of a tetrahedral mesh. Previous work has demonstrated that isotropic refinement is required if mesh quality is to be controlled effectively for arbitrary refinement levels. This is a serious limitation when directional flowfield features are present, leading to an inefficient distribution of grid points in the final mesh. In addition, truly anisotropic subdivision is almost impossible to realize for real problems on tetrahedral meshes.

A remedy for this drawback is to use hexahedral elements which can be subdivided anisotropically without mesh quality problems. Hexahedral meshes also yield more accurate flowfield solutions than their tetrahedral counterparts for the same number of edges. Our adaptation procedure uses an edge data structure that facilitates efficient subdivision by allowing individual edges to be marked for refinement or coarsening.

Hexahedral adaptation schemes generate hanging vertices when a hexahedron cannot be split into smaller hexahedra without continuously propagating the mesh refinement into regions where it is not desired. This problem is solved by using pyramids and prisms as buffers between refined and unrefined elements. These buffer elements are never subdivided, however. If an edge of a pyramid or a prism is marked for subdivision, then this element and its siblings coarsen back to their parent hexahedron and further refinement is performed directly on the hexahedral element. It is expected that this hexahedral mesh adaptation scheme will yield superior solutions with fewer mesh points than tetrahedral schemes.

HELICOPTER NOISE PREDICTIONS USING COMPUTATIONAL AEROACOUSTICS

Rupak Biswas and Roger C. Strawn (US Army AFDD)

High-performance helicopters and tiltrotors generate excessive noise both in forward flight and during takeoffs and landings. Accurate prediction of helicopter noise is essential to its control. Traditional acoustic analogy approaches cannot model the near-field nonlinear phenomena from high-speed rotor blades. CFD techniques are much better suited for these near-field nonlinearities but cannot efficiently propagate acoustic signals over long distances. The goal of this work is the accurate and efficient prediction of helicopter noise for a wide variety of flight regimes and rotor blade shapes using a hybrid CFD/Kirchhoff scheme.

We have used a solution-adaptive unstructured-grid Euler scheme with a stationary Kirchhoff method to model high-speed impulsive (HSI) noise for a hovering rotor. We have also developed a rotating Kirchhoff formulation to model HSI noise in forward flight. Both Kirchhoff methods have also been used to predict helicopter blade-vortex interaction (BVI) noise.

Up until now, these CFD/Kirchhoff techniques have been used to compute acoustic signals at a handful of far-field observer locations to compare with experimental microphone measurements. While these comparisons are useful for validation, they do not exploit the full capabilities of the new acoustic prediction methods. New tools are being developed for computing and analyzing helicopter noise in the far field. First, the Kirchhoff methods are being run at near ideal speedup on the SP2 so that several thousand acoustic signals can be computed in a few hours. Second, the predicted signals are postprocessed to obtain visual and auditory information about the resulting noise. A computer graphics program creates animated three-dimensional images of the acoustic pressure waves while an interactive audio system converts the computed digital acoustic signals into sound.

TENSOR METHODS FOR CONSTRAINED OPTIMIZATION

Dan Feng and Robert Schnabel (CU Boulder)

This research investigates tensor methods for nonlinear constrained optimization problems. These are general purpose methods especially intended for problems where the constraint gradient matrix is rank deficient or ill-conditioned at the solution. They are adapted from the standard successive quadratic programming (SQP) method by augmenting the linear model of the constraints with a simple second order term. The second order term is selected so that the model of the constraints interpolates constraint function values from several previous iterations, as well as the current constraint function value and gradients. Similarly to tensor methods for nonlinear equations, the tensor methods for constrained optimization require no more function and derivative evaluations, and hardly more storage or arithmetic per iteration, than the standard SQP methods. It is shown that the tensor methods are very efficient computationally on singular and well-conditioned nonlinear equality constrained optimization problems. A complete presentation of these methods is to appear in SIAM Journal on Optimization.

We analyze the local convergence properties of a version of the tensor method on problems where the constraint gradient matrix at the solution has rank deficiency one. The tensor model uses the same quadratic model of the Lagrangian as in the objective of the standard SQP model, and augments the linear model of the constraints in the SQP model by a rank one quadratic term. We show under certain conditions that the sequence of iterates generated by the tensor method based upon the tensor model converges locally and two-step Q-superlinearly to the solution with Q-order 3/2 on an interesting class of singular problems. In similar situation, we show that the standard SQP method converges only linearly with constant converging to 1/2. To the best of our knowledge, the result of this kind is the first for the SQP method. Hence, tensor methods have a theoretical advantage over the standard SQP method. Our analysis also confirms that the tensor method converges at least at the same rate as the standard SQP method on problems where the constraint gradient matrix at the solution has full rank. This research is presented in TR 95.10.

OPTIMALITY CONDITIONS FOR SINGULAR CONSTRAINED OPTIMIZATION

Richard Byrd (CU Boulder), Dan Feng and Robert Schnabel (CU Boulder)

Optimality conditions for singular constrained optimization are not well understood. When the constraint gradient matrix has rank deficiency at the solution, the Kuhn-Tucker conditions are usually not satisfied. We introduce several optimality conditions for singular constrained optimization problems for the case when the constraint gradient matrix at the solution is rank deficient. These conditions are a generalization of the traditional Kuhn-Tucker conditions for constrained optimization. These conditions have useful implications both for constructing algorithms and for analyzing methods for singular constrained optimization problems. This research is presented TR 95.03.

NEW ALGORITHMS FOR AERODYNAMIC DESIGN OPTIMIZATION

Dan Feng and Thomas H. Pulliam (NASA Ames)

A computational scheme is developed for solving a class of aerodynamic design problems that can be posed as nonlinear equality constrained optimization problems. The scheme treats the flow and design variables as independent variables, and solves the constrained optimization problem via reduced Hessian successive quadratic programming. It updates the design and flow variables simultaneously at each iteration and allows flow variables to be infeasible before convergence. The solution of an adjoint flow equation is never needed. In addition, a range space basis is chosen so that in a certain sense the ``cross term" ignored in reduced Hessian SQP methods is minimized. This scheme is extended to allow solution refining. In particular, we introduce a reduced Hessian refining technique that is critical for making a smooth transition of the Hessian information from coarse grids to fine grids. Through solution refining the efficiency and the robustness of the all-at-once reduced Hessian SQP scheme are significantly improved. This research is presented in TR 95.19 and 95.24.

ENERGY ESTIMATES FOR DIFFERENCE APPROXIMATIONS OF NONLINEAR CONSERVATION LAWS

Pelle Olsson

This project is a direct continuation of previous research projects carried out by the same author while at RIACS ("Summation by Parts, Projections, and Stability" and "Energy and Maximum Norm Estimates for Nonlinear Conservation Laws"). The ideas of TR 93.04 and TR 94.01 are transferred to semi-discrete systems. It turns out that the notion of entropy introduced by Lax is the natural extension of the linear L_2 stability concept to difference approximations of nonlinear conservation laws. We give several different characterizations of an entropy condition and show that they are all equivalent. The skew-symmetric form - which is an entropy condition - discussed in TR 94.01 is used to define a difference approximation of arbitrary order of accuracy. Using projections and summation by parts the resulting scheme is shown to satisfy an entropy condition, which in turn implies an energy estimate. The theory applies to initial-boundary value problems for symmetrizable systems of conservation laws in several space dimensions. The results are presented in TR 95.06.

HIGH-ORDER FINITE DIFFERENCE SHCEMES WITH SHARP SHOCK RESOLUTION FOR EULER EQUATIONS

Margot Gerritsen and Pelle Olsson

In previous research projects it has been demonstrated how to construct numerical schemes that support one- or two-point shocks for scalar conservation laws. This theory does not directly apply to the Euler equations, but following a similar approach we can construct a scalar viscosity that supports approximate one-point shocks. The viscosity is determined completely by the flow variables on either side of the shock.

The shock resolution can be significantly improved by the introduction of a small subgrid locally around the shock (steady as well as unsteady). The positioning of the subgrid is determined by a detection algorithm based on a multiscale wavelet analysis of the pressure grid function, which quickly and accurately locates potential shocks and spurious oscillations. It also supplies the information needed to compute the artificial viscosity terms. The detection algorithm is derived from a noise detection algorithm developed by Mallat and coworkers in signal analysis.

The resulting artificial viscosity can be incorporated into the difference scheme such that an entropy inequality holds.

PARALLEL PRECONDITIONER FOR CFD

Wei-Pai Tang

An effective linear solver is an essential part of sophisticated CFD code. Commonly, most of the CPU time of the computation is spent on the linear solver. When the size of the model grows and the difficulty of the simulation increases, the performance of the linear solver becomes crucial. If the computation is carried out in a high performance computer, new issues also arise.

The use of preconditioned Krylov space methods has been proven to be a competitive solution technique for wide ranges of large sparse matrix problems. It is acknowledged now that a high quality preconditioner is the key to the success. The objective of this project is to investigate parallel implementation issues when a high performance computer is used. This is a joint project with NAS Advanced Algorithm Division. It is an ongoing project. Fist phase of prototype testing is completed. A real parallel implementation is planed for next year.

RESEARCH IN AERODYNAMIC SHAPE OPTIMIZATION

James Reuther and Antony Jameson

Since the inception of CFD, researchers have sought not only accurate aerodynamic prediction methods for given configurations, but also design methods capable of creating new optimum configurations. Yet, while flow analysis can now be carried out over quite complex configurations using the Navier-Stokes equations with a high degree of confidence, direct CFD based design is still limited to very simple two-dimensional and three-dimensional configurations, usually without the inclusion of viscous effects. The CFD-based aerodynamic design methods that do exist can be grouped into two basic categories: inverse methods, and numerical optimization methods.

Inverse methods derive their name from the fact that they invert the goal of the flow analysis algorithm. Instead of obtaining the surface distribution of an aerodynamic quantity, such as pressure, for a given shape, they calculate the shape for a given surface distribution of an aerodynamic quantity. Most of these methods are based on potential flow techniques, and few of them have been extended to three-dimensions. The common trait of all inverse methods is their computational efficiency. Typically, transonic inverse methods require the equivalent of 2-10 complete flow solutions in order to render a complete design. Since obtaining a few solutions for simple two-dimensional and three-dimensional designs can be done in at most a few hours on modern computer systems, the computational cost of most inverse methods is considered to be minimal. Unfortunately, they suffer from many limitations and difficulties. Their most glaring limitation is that the objective is built directly into the design process and thus cannot be changed to an arbitrary or more appropriate objective function.

An alternative approach, which avoids some of the difficulties of inverse methods, but only at the price of heavy computational expense, is to use numerical optimization methods. The essence of these methods is very simple: a numerical optimization procedure is coupled directly to an existing CFD analysis algorithm. The numerical optimization procedure attempts to extremize a chosen aerodynamic measure of merit which is evaluated by the chosen CFD code. Most of these optimization procedures require gradient information in addition to evaluations of the objective function. Here, the gradient refers to changes in the objective function with respect to changes in the design variables. The simplest method of obtaining gradient information is by finite differences. In this technique, the gradient components are estimated by independently perturbing each design variable with a finite step, calculating the corresponding value of the objective function using CFD analysis, and forming the ratio of the differences. These methods are very versatile, allowing any reasonable aerodynamic quantity to be used as the objective function. They can be used to mimic an inverse method by minimizing the difference between target and actual pressure distributions, or may instead be used to maximize other aerodynamic quantities of merit such as L/D. Unfortunately, these finite difference numerical optimization methods, unlike the inverse methods, are computationally expensive because of the large number of flow solutions needed to determine the gradient information for a useful number of design variables. Tens of thousands of flow analyses would be required for a complete three-dimensional design. In our research, a new method is developed that avoids the limitations and difficulties of traditional inverse methods while retaining their inherent computational efficiency.

The method dramatically reduces the cost of aerodynamic optimization by replacing the expensive finite-difference method of calculating the required gradients with an adjoint variable formulation. After deriving the differential form of the adjoint equations and posing the correct boundary conditions based on the objective function, the resulting system is discretized and solved on the same mesh as that used for the flow solution. A significant economization is thus achieved by applying the same subroutines used for the flow solution to the solution of the adjoint equations. The resulting design process requires only one flow calculation and one adjoint calculation per gradient evaluation, as opposed to the hundreds required for a finite-difference gradient involving hundreds of design variables. In practice the computational cost of the new method is two orders of magnitude less then a conventional approach.

Considerable effort as been focused in the last two years to develop control theory based aerodynamic shape optimization methods. Especially in the last year emphasis has been placed on enhancements and developments that will allow these methods to be used as practical design tools as opposed to simple academic exercises. Such was the theme of our first paper, presented at the January 1995 Aerospace Sciences Meeting (AIAA paper 95-0123, also RIACS report 95.01), which showed how the basic techniques that we had developed in the previous year for two-dimensional airfoil configurations could easily be extended to three-dimensional wing and wing-body configurations.

One of the fallouts from this presentation was its stimulation of interest from the Beechcraft Aircraft Division of Raytheon, Inc. Raytheon entered into a cooperative agreement with NASA Ames Research Center to explore the usefulness of the adjoint based design optimization methods we have been developing. In the middle of March, representatives from Raytheon were on site at Ames to test the new design techniques on a new transonic wing that they were developing for an all-new business jet they proposed to build and market. Since at that time we were able to treat only the design of wing-body configurations and their design involved a wing-body-nacelle configuration, some very imaginative design strategies were developed in order to permit our methods to be applicable to their very complex real world problem. By the beginning of May, a new wing had been designed using the adjoint based method and validated computationally. This one-month design of a new transonic wing compares with the usual development time of more than a year for traditional methods. Raytheon has since wind tunnel tested the new wing design to confirm its predicted performance, and launched the design for production. They took 51 orders for the new airplane on the day they announced the design. Furthermore, Raytheon has been so impressed by the capability of adjoint based design methods that they are now incorporating them into their own aircraft design environments. A paper authored by both NASA and Raytheon personnel that presents the basic design strategy and its outcome will be presented at the Aerospace Sciences Meeting set for January 1996 (AIAA paper 96-0554, also RIACS 96.03).

Another group that has taken a keen interest in our research is the NASA High Speed Research Program (HSR) group. In their effort to create economically viable supersonic transport configurations for the next century they are investigating the use of aerodynamic shape optimization to improve performance. Both the traditional as well as adjoint based design methods created by our group at Ames have been tested by the HSR community. During the summer our group supported many wind tunnel tests at NASA Langley Research Center in Virginia to validate the performance gains predicted by our aerodynamic shape optimization capability. A paper that gives an example of the capabilities of this emerging technology for supersonic design was presented at the American Society of Mechanical Engineers annual winter meeting in November 1995 (also RIACS 95.14).

The experience of working with Raytheon and the HSR community has forced us to consider many aspects of aerodynamic shape design that are often neglected from the purely academic standpoint. One of the many new areas of research that will develop as adjoint based design becomes established is the study of possible aerodynamic design space parameterizations. We have made a first look into this area through a paper that explores the use of both Hicks-Henne functions and B-spline control points as design variables. The paper was presented at the Sixth International Symposium on Computational Fluid Dynamics conference in September 1995 (also RIACS report 95.13). In another effort to enhance our capabilities, the methodology was extended for the treatment of multiblock structured three-dimensional meshes. To accomplish this task we devised three new elements: a multiblock flow solver, a multiblock adjoint solver, and a multiblock mesh perturbation method. Initial results indicate that as of the end of the year the entire system is now operational and permits complete aircraft configurations to be treated as part of the design process. This method can now treat the complex wing-body-nacelle design needed by Raytheon without resorting to difficult iterative strategies. A paper to be presented in the forthcoming 34th Aerospace Sciences Meeting (AIAA paper 96-0094, also RIACS report 96.02) presents this new multiblock design method.

In spite of the major accomplishments achieved this year, much more research will be required in order to harness the true potential of adjoint based aerodynamic design. Nevertheless, the fact that we have been able to make significant improvements on both existing and future designs using our new technologies has demonstrated beyond doubt the great value of adjoint based aerodynamic design. It is hoped that with all of these advances, the greater aeronautical science community will in the future adopt these new ideas into their production design environments. Certainly if the work in conjunction with Raytheon is any indication, this is already taking place.

ADAPTIVE REFINEMENT OF COMPOSITE CURVILINEAR GRIDS

Steven Suhr

A software system is being designed and implemented to manage the adaptive refinement of composite curvilinear grids for the approximate solution of time-dependent partial differential equations. With the simplifying assumption that the spatial domain has fixed boundaries, an initial grid is constructed using a fixed set of overlapping grids, which collectively conform to the boundaries and cover the domain. Refinement grids, aligned with the original base grids, are added to maintain accuracy as the solution evolves. This approach organizes the grids into a geometrically nested tree of connected components, and it explores the use of curvilinear stairstep refinement grids.

The programming language Vorpal, currently under development, will be used in the implementation of this system, taking advantage of Vorpal's support for data structures, abstract data types, structured external files, and modular program structure. An important milestone in the transformation of Vorpal from a collection of useful concepts into a unified preliminary design will be the translation of adaptive grid code being implemented at Stanford by others for model problems in two space dimensions, from C into Vorpal. When an implementation of Vorpal exists, the adaptive grid system will be extended from two to three space dimensions, and a version which can be applied more readily to realistic and diverse problems will be created. As the adaptive grid system evolves and grows, the anticipated future support in Vorpal for concurrency should also be useful.

C. HIGH PERFORMANCE NETWORKS

BAY AREA GIGABIT NETWORK TESTBED

Marjory J. Johnson and Jeffrey N. Townsend

The BAGNet testbed, sponsored by Pacific Bell's CalREN (California Research and Education) program, has been in existence for almost two years. This testbed involves fifteen research, educational, governmental, and industrial organizations within the Bay Area. M. Johnson, Bill Johnston of LBL, and Dan Swinehart of Xerox PARC are jointly coordinating the testbed project.

In this section we discuss general testbed activities over the past year; accomplishments regarding specific testbed-oriented RIACS projects are presented in separate sections.

The basic infrastructure for the testbed remains the same as last year. BAGNet is an IP over ATM testbed, using OC-3c (155 Mbps) links. A mesh of permanent virtual channels provides connections between all possible pairs of testbed hosts. In addition, each host has one point-to-multipoint PVC for outgoing multicast transmission to all other testbed hosts.

BAGNet has been an ambitious project, both because of the immature level of ATM technology when the testbed originated, and because of the large number of testbed participants and the heterogeneity of equipment involved. Major testbed accomplishments over the past year include identification and resolution of many issues that must be addressed when building a relatively large-scale IP-over-ATM network, development of a methodology to maintain up-to-date configuration status, analysis of performance degradation caused by bandwidth-policing policies, compilation of comprehensive performance statistics obtained by using various performance tools (e.g., Netperf), and implementation of multicast capabilities by using point-to-multipoint PVCs. We have not been able to experiment with specifying and delivering quality-of-service guarantees, since switched virtual channels are not yet available for us to use. More detailed information about the above activities can be obtained from the World Wide Web, via links from the general BAGNet home page (http://george.lbl.gov/BAGNet.html).

Some of the testbed sites are now broadcasting seminars on a regular basis over BAGNet. The quality of reception varies, depending on individual workstation capabilities. Clearly, the bottleneck for performance is not the network, but rather workstation architecture and protocol issues.

Late in the year we initiated a collaboration with Bellcore to capture data for traffic analysis, so that we can see what a data stream for a real ATM application looks like. Bellcore is interested in general traffic analysis to help them understand how to manage data transmission for commercial-service offerings; several of the individual sites plan to use the collected data to model specific applications. We are also interested in observing and analyzing interactions between applications that are sharing the testbed.

At RIACS we are developing and testing protocols for file transfer that will utilize the available high bandwidth of the testbed. We are also developing an environment to support collaborative scientific work. Other testbed projects within NASA Ames include video on demand (VOD) and remote access to NASA wind tunnels.

COLLABORATIVE SCIENCE

Marjory J. Johnson and Jeffrey N. Townsend

The goal of this project is to develop and evaluate an environment to support real-time collaborative scientific work, using high-speed ATM networks as the communications medium. Project collaborators include earth scientists at Ames and at the University of Arizona and networking researchers at Sandia National Laboratories - Livermore.

We have established a collaborative environment that incorporates tools developed by Sandia. This year we have enhanced the infrastructure for our project by adding video capabilities to our BAGNet workstation, enabling a virtual presence with remote colleagues.

We have begun experiments over BAGNet between RIACS and Sandia. During these experiments, we quickly found that the Sparc 2 architecture of our BAGNet workstation is inadequate to take full advantage of the bandwidth capabilities offered by BAGNet. We are currently working to upgrade our workstation to a Sparc 20 architecture.

Recently Bellcore requested BAGNet participation in a project to record actual ATM application network traffic. We conducted a typical collaborative work session while Bellcore was recording BAGNet testbed traffic. In the future we will analyze this traffic data, to obtain an understanding of communication patterns for collaborative applications.

Future plans also include working with earth scientists at NASA Ames to evaluate the tools we have incorporated in our collaborative environment. With their assistance we will develop end-user applications, e.g., interactive data analysis and remote training, that incorporate this environment and that are enabled by it.

HIGH-RATE DATA TRANSFER

Marjory J. Johnson and Jeffrey N. Townsend

The standard file-transfer protocol, ftp, yields throughput rates that are disappointingly low, relative to the raw bandwidth that is available with high-performance fiber-optic networks such as the BAGNet testbed. Disk I/O is clearly a primary bottleneck, and sender/receiver interactions to ensure error-free transmission provide another source of significant overhead. Such issues indicate that new data-transfer protocols must be designed to utilize the bandwidth that is available with emerging network technologies.

Our application focus is the transfer of large image files. Since many image-transfer applications can tolerate a low level of transmission errors, we are basing our protocol development on UDP rather than TCP. Our goal is to develop a data-transfer protocol that maximizes network throughput over ATM networks, while keeping transmission errors manageable. Of course, the level of transmission errors that is considered acceptable is application dependent.

We are currently experimenting with several techniques for data transfer, all of which attempt to keep the transmission pipe full. We are using multiple data streams, so that disk I/O, etc. can be overlapped with data transfer. Transmission errors are controlled via a low level of synchronization of sender/receiver activities.

We are conducting our experiments on a variety of workstations located at three BAGNet sites: NASA Ames, Sandia National Laboratories, and Lawrence Livermore National Laboratory. Preliminary results are contained in the paper, ``Achieving High Throughput for Data Transfer over ATM Networks," which has been accepted for presentation at the International Communications Conference, ICC '96.

Early results validate our approach. We are able to limit transmission errors to two to three percent, while achieving throughput rates that are several times higher than ftp rates.

In the future we will continue to refine our protocols for data transfer. Then we will evaluate the performance of our protocol as part of an end-user application that involves image transfer.

D. LEARNING SYSTEMS

SUPER-RESOLUTION/SURFACE MODELING

Peter Cheeseman

Proposed Capability: An algorithm that takes multiple images of the same area/object and combines the information to give a much higher resolution image, and identification of ground truth in the image.

Benefits: More science from images from Viking, Voyager, Galileo, LandSat, Tiros etc., through improved resolution. Also, loss less data compression and change detection for Earth observation data. An 8:1 increase in resolution has already been demonstrated.

User Interest: Preliminary discussions with LORAL (EOS-DIS contractor) have identified potential areas for collaboration. A demonstration project with University of Colorado at Boulder [Bill Emory] has developed software for earth science applications. Presentation to planetary scientists has resulted in collaborations to produce higher resolution images of sites of considerable scientific interest.

Industry and Academic Involvement: Collaboration and co-funding of development of our algorithms for earth science applications at the University of Colorado at Boulder.

Commercial Use: A number of companies have requested our super-resolution code.

Beginning in late 1991, the Bayes group at ARC began a project to develop the theory and practice of multiple image data combination. Multiple images taken from similar locations and under similar lighting conditions contain similar - but not identical - information. Slight differences in instrument orientation and position produce mismatches between the pixel grids of different images. These mismatches ensure that any point on the ground is sampled differently in each image. The surface modeling project is designed to exploit these differences to build a super-resolved composite image that uses all the information from the separate images.

The basic theory behind our approach is that of inverse graphics. That is, if we knew what the ground is like, the lighting conditions, and the camera orientation and characteristics, etc., then we could predict what the camera would see (an image). This is the standard computer graphics problem. However, we have the inverse problem–we know what the images are, and we want to find the most probable ground truth (surface) that would have generated them, assuming we know the lighting conditions and camera characteristics. The most important (and difficult) part of this process is recovering the camera orientation and position for each image. To do this, we must register all the images with respect to each other to an accuracy of a small fraction of a pixel; this registration tells us how an image maps onto the ground truth model we are building. Our initial ground model is formed by letting each pixel "vote" on what the corresponding ground position should be depending on how much that ground position contributed to that pixel. This initial ground model is then used to project what each image should be (i.e., predict each pixel value). The differences between the predicted pixel value and the observed value is used to update the ground model until it cannot be further improved. This procedure produces an increase in both spatial resolution and gray-scale resolution.

The above procedure for super-resolution surface reconstruction from multiple images could be implemented with standard statistical (maximum likelihood) methods. The distinctly Bayesian contribution is to include neighbor correlation into the ground pixels model. In other words, in the surface reconstruction, a ground pixel is not just a function of the corresponding projected pixels, but also depends on its neighbor values. This correlation addition makes our approach much less sensitive to noise in the pixel values, and produces more realistic reconstruction. Test results with 8:1 resolution enhancement, and Viking Orbiter imagery with 2:1 to 4:1 enhancements are attached.

Fiscal Year 1995

Progress In FY95 we have refined our registration algorithm so that we can achieve registration to at least within 1/10th. of a pixel. This very accurate registration has allowed us to obtain 8:1 improvement in spatial resolution. In the process of developing this accurate registration process, we had to invent a new procedure for function minimization from generated function values that maximized the use of information in these very expensive function values. This new procedure has applications well beyond our particular problem. We have also increased the speed and reliability of the code, and tested it on scientifically important test cases on Mars. The original super-resolution Lisp version has been rewritten in C, as requested by several of our users.

Our test cases to date have effectively been 2-D surfaces (small patches of Mars or Earth). However, the Voyager imagery of the moons of Jupiter and Saturn, required us to extend the theory to projections onto spherical surfaces. We are currently testing this extended theory, and expect to get improved resolution images of some of these moons that exceeds the best currently available. We are also developing more extensive 3-D super-resolution surface reconstruction for irregular shaped objects such as Gaspra (an asteroid) from images obtained by the Galileo fly-by.

In 1994 and again in 1995, we began a collaborative research arrangement with the University of Colorado at Boulder to apply the super-resolution techniques to earth observing satellite images. This collaboration involves developing a 3-D surface model of selected earth test sites, using a USGS elevation model as the starting point. We have developed techniques for taking satellite earth observation data, particularly AVHRR data available daily, and integrating this information into the high resolution 3-D grid. Even though the AVHRR pixel size is 1 kilometer square, at best, we expect to get roughly an order of magnitude improvement in resolution by integrating many images. This approach allows us to use other satellite data, such as GOES, OLS, SPOT and LandSat, depending on availability.

Fiscal Year 1996 Plan

Our current algorithm is limited to flat areas with nearly identical viewing angle and lighting conditions. We plan to extend the theory and implementation surface reconstruction for spherical planetary surfaces, and apply the resulting algorithms to all Voyager planetary images of scientific interest. We will also continue the development of the theory and implementation of arbitrary 3-D surface models, with any combination of lighting and viewing conditions. This extension will require modeling of shadows and atmospheric conditions. These extensions will allow us to produce super-resolved images (3-D models) of objects such as asteroids and close-ups of earth's surface.

These extensions require that we extend the theory to integrate information across different spectral bands. We expect to be able to do this by allowing correlations between the bands that vary slowly across the image. For earth images we must also develop atmospheric models to remove the effects of clouds and haze from individual images. We expect that information from the half-hourly GOES images will help separate the rapidly changing cloud effects from the relatively constant earth background. We also plan to develop "classification models" where each ground grid point is assigned a percentage membership of all the known classes of ground cover (e.g., soil, grass, trees, etc.). The spectral characteristics of each class is known, so it is possible to calculate, for each grid point, what the expected spectral characteristics of that point should be. The observed characteristics allow us to update the classification vector at each point, including the assumption that the classification vectors for adjacent points are highly correlated. This kind of classification model will be extremely useful for quantitative earth science studies.

Primary Users: EOS-DIS, earth scientists, planetary scientists. Potential Users: Anyone trying to get more information out of a set of images—e.g., high resolution television; improved microscope and telescope resolution; higher fax resolution; etc. Delivery Year of Major Capability: 1996

KNOWLEDGE COMPILATION AND LEARNING FOR EFFICIENT GENERAL PROBLEM-SOLVING

Barney Pell

This research investigates methods to increase the generality of problem-solving systems by combining general-purpose knowledge with details of a specific problem to be solved in order to produce an efficient special-purpose system to solve that problem efficiently. This separation of general knowledge from specifics of a given problem provides several advantages. It facilitates transfer of knowledge from one problem domain to another, enables a user to specify the problem in high-level language without confusing the problem with its implementation, and the concise representation resulting from such a problem-formulation enables the machine to learn from experience across a variety of domains. Current work centers on extending the author's Ph.D. thesis research, which focused on games (like chess and checkers) to problems of industrial importance. Such problems include information retrieval and automatic scheduling.

INDUCTION FOR HEURISTIC INFORMATION RETRIEVAL

Barney Pell, Catherine Baudin (NASA/Recom) and
Smadar Kedar (ILS, Northwestern University)

This research involves developing methods for using supervised learning techniques to aid information retrieval (IR). A general problem within IR is that a user's QUERY may have no direct match to any indexed document, but it would be desirable for the system to return something related to that query anyway. The framework we are using involves having a knowledge engineer design a set of RETRIEVAL HEURISTICS. When the indexing system has no direct match to a user's query, the heuristics operate on the query to generate an extended query set, and the cases matching these extended queries are offered to the user. We are exploring how a system can use FEEDBACK from users to improve its retrieval performance.

AUTONOMOUS SYSTEM ARCHITECTURES

Barney Pell, Doug Bernard (JPL), Steve Chien (JPL), Erann Gat (JPL),
Nicola Muscettola (Recom), Pandu Nayak (Recom),
Mike Wagner (Air Force Phillips Lab& Code IC), Brian Williams (Recom)

Two of the primary cost drivers for current space programs are mission operations (ground-based sequencing and spacecraft heath maintenance) and flight software development and integration. The cost of mission operations can be reduced through automation, and software development costs can be reduced by reuse of integrated autonomous software across missions and across spacecraft. The cost-effective design and reuse of autonomous systems software is challenging because of problems of overwhelming detail, system-wide interactions, and integration of functionalities. This research investigates the development of autonomous spacecraft architectures. These architectures raise the level of abstraction that the spacecraft designer must consider, and they support modular software design. They also facilitate the development and integration of autonomous components for planning, anomaly detection, diagnosis, resource allocation and robust execution.

III. TECHNICAL REPORTS

TR 95.01

AERODYNAMIC SHAPE OPTIMIZATION OF WING AND WING-BODY CONFIGURATIONS USING CONTROL THEORY

Antony Jameson, James Reuther
January 1995 (31 pages)

Proceedings of the the 33rd Aerospace Sciences Meeting and Exhibit, January, Reno, NV. AIAA paper number 95-0123.

This paper describes the implementation of optimization techniques based on control theory for wing and wing-body design. In previous studies it was shown that control theory could be used to devise an effective optimization procedure for airfoils and wings in which the shape and the surrounding body-fitted mesh are both gene Recently, the method has been implemented for both potential flows and flows governed by the Euler equations using an alternative formulation which employs numerically generated grids, so that it can more easily be extended to treat general configurations. Here results are presented both for the optimization of a swept wing using an analytic mapping, and for the optimization of wing and wing-body configurations using a general mesh.

TR 95.02

AN ACCURACY TEST OF A CARTESIAN GRID METHOD FOR STEADY FLOW IN COMPLEX GEOMETRIES

Marsha Berger (Courant Institute, New York) and
John Melton (NASA Ames Research Center)
February 1995 (10 pages)

Proceedings of the 5th International Conference on Hyperbolic Problems, Stonybrook, NY, June 1995

We describe a method to solve the inviscid compressible Euler equations using non-body-fitted Cartesian grids. We briefly survey the geometric procedures needed for the "grid generation" in this approach and show some Cartesian grids for three dimensional configurations. The major source of error in these methods is at the irregular boundary cells. This has motivated us to study various improvements in the numerical scheme (a Godunov method) on a model problem with smooth flow in two space dimensions.

LIMITED NUMBER OF 95.02 AVAILABLE DUE TO REPRODUCTION COSTS

TR 95.03

ON OPTIMALITY CONDITIONS FOR SINGULAR CONSTRAINED OPTIMIZATION

Richard H. Byrd (University of Colorado, Boulder), Dan Feng and Robert B. Schnabel (University of Colorado, Boulder)
March 1995 (17 pages)

Optimality conditions for singular constrained optimization are not well understood. When the constraint gradient matrix is rank deficient at the solution, the Kuhn-Tucker conditions may not be satisfied. This paper introduces and discusses several optimality conditions for singular constrained optimization problems for the case when the constraint gradient matrix at the solution is rank deficient. These conditions are a generalization of the traditional Kuhn-Tucker conditions for constrained optimization. A class of ``nice" singular constrained problems is identified and is shown to satisfy these conditions. We also give examples of singular constrained problems that do not satisfy these conditions, and point out fundamental differences between these problems and the nice problems. These conditions have useful implications both for constructing algorithms and for analyzing methods for singular constrained optimization problems.

TR 95.04

A PORTABLE MPI-BASED PARALLEL VECTOR TEMPLATE LIBRARY

Thomas J. Sheffler
February 1995 (26 pages)

This paper discusses the design and implementation of a poly morphic collection library for distributed address-space parallel computers. The library provides a data-parallel programming model for C++ by providing three main components: a single generic collection class, generic algorithms over collections, and generic algebraic combining functions. Collection elements are the fourth component of a program written using the library and may be either of the built-in types of C or of user-defined types. Many ideas are borrowed from the Standard Template Library (STL) of C++, although a restricted programming model is proposed because of the distributed address-space memory model assumed. Whereas the STL provides standard collections and implementations of algorithms for uniprocessors, this paper advocates standardizing interfaces that may be customized for different parallel computers. Just as the STL attempts to increase programmer productivity through code reuse, a similar standard for parallel computers could provide programmers with a standard set of algorithms portable across many different architectures. The efficacy of this approach is verified by examining performance data collected from an initial implementation of the library running on an IBM SP-2 and an Intel Paragon.

TR 95.05

HELICOPTER NOISE PREDICTIONS USING KIRCHHOFF METHODS

Rupak Biswas, Roger C. Strawn (US Army AFDD) and Anastasios S. Lyrintzis, (Purdue University)
February 1995 (14 pages)

This paper presents two methods for predicting the noise from helicopter rotors in forward flight. Aerodynamic and acoustic solutions in the near field are computed with a finite-difference solver for the Euler equations. Two different Kirchhoff acoustics methods are then used to propagate the acoustic signals to the far field in a computationally-efficient manner. One of the methods uses a Kirchhoff surface that rotates with the rotor blades. The other uses a nonrotating Kirchhoff surface. Results from both methods are compared to experimental data for both high-speed impulsive noise and blade-vortex interaction noise. Agreement between experimental data and computational results is excellent for both cases. The rotating and nonrotating Kirchhoff methods are also compared for accuracy and efficiency. Both offer high accuracy with reasonable computer resource requirements. The Kirchhoff integrations efficiently extend the near-field finite-difference results to predict the far field helicopter noise.

TR 95.06

SUMMATION BY PARTS, PROJECTIONS, AND STABILITY III

Pelle Olsson
April 1995 (32 pages)

A new way of characterizing an entropy pair (η, q) by means of Euler's inhomogeneous differential equation is proposed. This characterization is shown to be equivalent to existing definitions. It is then demonstrated how to obtain entropy-conservative finite difference schemes of arbitrary order of accuracy for initial-boundary value problems of nonlinear symmetrizable systems of conservation laws. No smallness assumption on the oscillation of the solution is necessary. The entropy conservation follows by applying the energy method to a difference scheme that is based on a certain splitting of the flux derivative, where Euler's inhomogeneous differential equation again plays a crucial role. The entropy functions for the continuous and the discrete problems are identical. Finally, using the entropy conservation it is shown how to construct a generalized energy estimate. In fact, the energy estimates show that the difference schemes are strictly stable. The analysis is carried out in one space dimension, but can be generalized to multidimensional domains.

TR 95.07

A FAST POISSON SOLVER FOR UNSTEADY INCOMPRESSIBLE NAVER-STOKES EQUATIONS ON THE HALF-STAGGERED GRID

G. H. Golub (Stanford University), L.C. Huang (Laboratory for Scientific and Engineering Computing), H. Simon (Silicon Graphics), W. -P. Tang
April 1995 (17 pages)

In this paper, a fast Poisson solver for unsteady, incompressible Navier-Stokes equations with finite difference methods on the nonuniform, half-staggered grid is presented. To achieve this, new algorithms for diagonalizing a semi-definite pair are developed. Our fast solver can also be extended to the three dimensional case. The motivation and related issues in using this second kind of staggered grid are also discussed. Numerical testing has indicated the effectiveness of this algorithm.

TR 95.08

LEARNING IN NETWORKS

Wray L. Buntine
April 1995

Intelligent systems require software incorporating probabilistic reasoning, and often times learning. Networks provide a framework and methodology for creating this kind of software. This paper introduces network models based on chain graphs with deterministic nodes. Chain graphs are defined as a hierarchical combination of Bayesian and Markov networks. To model learning, plates on chain graphs are introduced to model independent samples. The paper concludes by discussing various operations that can be performed on chain graphs with plates as a simplification process or to generate learning algorithms.

TR 95.09

A SOLUTION ADAPTIVE STRUCTURED/UNSTRUCTURED OVERSET GRID FLOW SOLVER WITH APPLICATIONS TO HELICOPTER ROTOR FLOWS

Earl P. N. Duque, US Army AFDD, Rupak Biswas, Roger C. Strawn, US Army AFDD
April 1995 (10 pages)

This paper summarizes a method that solves both the three dimensional thin-layer Navier-Stokes equations and the Euler equations using overset structured and solution adaptive unstructured grids with applications to helicopter rotor flowfields. The overset structured grids use an implicit finite-difference method to solve the thin-layer Navier-Stokes/Euler equations while the unstructured grid uses an explicit finite-volume method to solve the Euler equations. Solutions on a helicopter rotor in hover show the ability to accurately convect the rotor wake. However, isotropic subdivision of the tetrahedral mesh rapidly increases the overall problem size.

TR 95.10

LOCAL CONVERGENCE ANALYSIS OF TENSOR AND SQP METHODS FOR SINGULAR CONSTRAINED OPTIMIZATION

Dan Feng, Robert B. Schnabel, University of Colorado, Boulder
June 1995 (33 pages)
Submitted to: SIAM Journal on Optimization

Tensor methods for nonlinear constrained optimization are adapted from the standard successive quadratic programming (SQP) method by augmenting the linear model of the constraints with a simple second order term. These methods have been shown to be efficient and robust computationally, especially on problems where the constraint gradient matrix at the solution has a small rank deficiency. This paper analyzes the local convergence properties of a version of tensor methods, on problems where the constraint gradient matrix at the solution has rank deficiency one. The tensor model that is analyzed uses the same quadratic model of the Lagrangian as in the objective of the standard SQP model, and augments the linear model of the constraints in the SQP model by a rank one quadratic term. We show that under certain conditions the sequence of iterates generated by the tensor method converges locally and two-step Q-superlinearly to the solution with Q-order 3/2 on an interesting class of singular problems. In a similar situation, we show that the standard SQP method converges linearly with constant converging to 1/2. Hence, tensor methods have a theoretical advantage over the standard SQP method on certain singular problems. Our analysis also confirms that the tensor method converges at least the same rate as the standard SQP method on problems where the constraint gradient matrix at the solution has full rank.

TR 95.11

DATA STRUCTURES AND IMPLEMENTATION OF COMPOSITE ADAPTIVE GRID METHODS

Joseph Oliger, Xiaolei Zhu
June 1995 (34 pages)
Submitted to: Journal of Applied and Computational Mathematics

Composite adaptive grid (CAG) methods for solving hyperbolic partial differential equations (PDE's) are discussed in this paper. Composite grids are used at the coarsest level to deal with domains with complicated geometries. Stair step grids developed for composite grids are also used in the finer levels to cover refined regions. New data structures for stair step grids are proposed and implemented in two and three space dimensions. Good efficiency is demonstrated using the numerical examples.

TR 95.12

ON SUPER-STABLE IMPLICIT METHODS AND TIME-MARCHING APPROACHES

Peter K. Sweby (University of Reading), Helen C. Yee (NASA Ames Research Center)
July 1995 (37 pages)
Submitted to: International J. of Bifurcation and Chaos
Proceedings of the Conference on Numerical Methods for Euler and Navier-Stokes Equations, Sept. 14-16, 1995.

The objective is to utilize dynamical system theory to illustrate some basic ideas of how super-stable implicit methods can influence the reliability, efficiency, stability, and convergence properties of time-dependent approaches for obtaining steady-state numerical solutions. We show how these recent advances can provide more reliable criteria and guidelines for the usage of commonly used implicit linear multistep methods in CFD. The knowledge gained may minimize the tuning of computational parameters and improve efficiency of the numerical procedure for practical CFD computations. Of particular interest are examples of the nonlinear effect caused by grid adaptation and super-stable implicit TVD schemes on the over all performance of the numerical procedure. A guideline to avoid spurious steady-state numerical solutions and asymptotes is included.

LIMITED NUMBER OF 95.12 AVAILABLE DUE TO REPRODUCTION COSTS

TR 95.13

A COMPARISON OF DESIGN VARIABLES FOR CONTROL THEORY BASED AIRFOIL OPTIMIZATION

Antony Jameson, James Reuther
July 1995

This paper describes the implementation of optimization techniques based on control theory for airfoil design. In our previous work in the area it was shown that control theory could be employed to devise effective optimization procedures for two-dimensional profiles by using either the potential flow or the Euler equations with either a conformal mapping or a general coordinate system. We have also explored three-dimensional extensions of these formulations recently The goal of our present work is to demonstrate the versatility of the control theory approach by designing airfoils using both Hicks-Henne functions and B-spline control points as design variables. The research also demonstrates that the parameterization of the design space is an open question in aerodynamic design.

TR 95.14

SUPERSONIC WING AND WING-BODY SHAPE OPTIMIZATION USING AN ADJOINT FORMULATION

Antony Jameson and James Reuther
August 1995 (8 pages)
Proceedings of the Forum on CFD for Design and Optimization, International Mechanical Engineering Conference and Exposition (IMECE 95) November 12-17, 1995, San Francisco, CA.

This paper describes the implementation of optimization techniques based on control theory for wing and wing-body design of supersonic configurations. The work represents an extension of our earlier research in which control theory is used to devise a design procedure that significantly reduces the computational cost by employing an adjoint equation. In previous studies it was shown that control theory could be used to devise transonic design methods for airfoils and wings in which the shape and the surrounding body-fitted mesh are both generated analytically, and the control is the mapping function. The method has also been implemented for both transonic potential flows and transonic flows governed by the Euler equations using an alternative formulation which employs numerically generated grids, so that it can treat more general configurations. Here results are presented for three-dimensional design cases subject to supersonic flows governed by the Euler equation.

TR 95.15

PARALLEL IMPLEMENTATION OF AN ADAPTIVE SCHEME FOR 3D UNSTRUCTURED GRIDS ON A SHARED-MEMORY MULTIPROCESSOR

Rupak Biswas and Leonardo Dagum (Silicon Graphics, Inc.)
August 1995 (8 pages)

Dynamic mesh adaption on unstructured grids is a powerful tool for computing unsteady flows that require local grid modifications in order to efficiently resolve solution features. For such flows, the coarsening/refinement step must be completed every few time steps, so its efficiency must be comparable to that of the flow solver. For this work, we consider an edge-based adaption scheme that has shown good single processor performance on a Cray Y-MP and C90. We report on our experience porting this code to an SGI Power Challenge system and parallelizing it for shared-memory symmetric multiprocessor architectures.

TR 95.16

AN ASPECT RATIO BOUND FOR TRIANGULATING a d--GRID CUT BY A HYPERPLANE

Scott A. Mitchell (Sandia National Laboratories), Stephen Vavasis
August 1995 (20 pages)

We consider the problem of triangulating a d-dimensional uniform grid of d-cubes that is cut by a k-dimensional affine subspace. The goal is to obtain a triangulation with bounded aspect ratio. To achieve this goal, we allow some of the box faces near the affine subspace to be displaced. This problem has applications to finite element mesh generation. For general d and k, the bound on aspect ratio that we attain is double-exponential in d. For the important special case of d=3, the aspect ratio bound is small enough that the technique is useful in practice.

TR 95.17

SEARCHING FOR PATTERNS IN REMOTE SENSING IMAGE DATABASES USING NEURAL NETWORKS

Justin D.Paola, Robert A. Schowengerdt (University of Arizona, Tucson)
Proceedings of the 15th Annual IEEE International Geoscience and Remote Sensing Symposium, Florence, Italy, July 20-14, 1995, pp. 443-445.
August 1995 (10 pages)

We have investigated a method, based on a successful neural network multispectral image classification system, of searching for single patterns in remote sensing databases. While defining the pattern to search for and the feature to be used for that search (spectral, spatial, temporal, etc.) is challenging, a more difficult task is selecting competing patterns to train against the desired pattern. Schemes for competing pattern selection, including random selection and human interpreted selection, are discussed in the context of an example detection of dense urban areas in LandSat Thematic Mapper imagery. When applying the search to multiple images, a simple normalization method can alleviate the problem of inconsistent image calibration. Another potential problem, that of highly compressed data, was found to have a minimal effect on the ability to detect the desired pattern. The neural network algorithm has been implemented using the PVM (Parallel Virtual Machine) library and nearly-optimal speedups have been obtained that help alleviate the long process of searching through imagery.

TR 95.18

THE EFFECT OF LOSSY IMAGE COMPRESSION ON IMAGE CLASSIFICATION

Justin D.Paola, Robert A. Schowengerdt (University of Arizona, Tucson)
Proceedings of the 15th Annual IEEE International Geoscience and Remote Sensing Symposium, Florence, Italy, July 20-14, 1995, pp. 118-120.
August 1995 (12 pages)

We have classified four different images, under various levels of JPEG compression using the following classification algorithms: minimum-distance, maximum-likelihood, and neural network. The training site accuracy and percent difference from the original classification were tabulated for each image compression level, with maximum-likelihood showing the poorest results. In general, as compression ratio increased, the classification retained its overall appearance, but much of the pixel-to-pixel detail was eliminated. We also examined the effect of compression on spatial pattern detection using a neural network.

TR 95.19

AL ALL-AT-ONCE REDUCED HESSIAN SQP SCHEME FOR AERODYNAMICS DESIGN OPTIMIZATION

Dan Feng and Thomas H. Pulliam (NASA Ames Research Center)
16 pages, Oct 1995

This paper introduces a computational scheme for solving a class of aerodynamic design problems that can be posed nonlinear equality constrained optimization. The scheme treats the flow and design variables as independent variables, and solves the constrained optimization problem via reduced Hessian successive quadratic programming. It updates the design and flow variables simultaneously at each iteration and allows flow variables to be infeasible before convergence. The solution of an adjoint flow equation is never needed. In addition, a range space basis is chosen so that in certain sense the ``cross term" ignored in reduced Hessian SQP methods is minimized. Numerical results on a nozzle design using quasi-one-dimensional Euler equations show that this scheme is computationally efficient and robust. The computational cost of a typical nozzle design is only a fraction more than that of the corresponding analysis flow calculation. Superlinear convergence is also observed, which agrees with the theoretical properties of this scheme. All optimal solutions are obtained by starting far away from the solution.

TR 95.20

SPECTAL ORDERING TECHNIQUES FOR INCOMPLETE LU PRECONDITIONERS FOR CG METHODS

Simon S. Clift (Queens University, Canada), Horst D. Simon (Silicon Graphics, Inc.) and Wei-Pai Tang
September 1995 (23 pages)

The effectiveness of an incomplete LU (ILU) factorization as a preconditioner for the conjugate gradient method can be highly dependant on the ordering of the matrix rows during its creation. Detailed justification for two heuristics commonly used in matrix ordering for anisotropic problems is given. The bandwidth reduction, and weak connection following heuristics are implemented through an ordering method based on eigenvector computations. This spectral ordering is shown to be a good representation of the heuristics. Analysis and test cases in two and three dimensional diffusion problems demonstrate when ordering is important, and when an ILU decomposition will be ordering insensitive. The applicability of the heuristics is thus evaluated and placed on a more rigorous footing.

TR 95.21

EXPERIENCES WITH THE BAY AREA GIGABIT NETWORK TESTBED

Marjory J. Johnson
Proceedings of the Fifth IEEE Workshop on Future Trends of Distributed Computing Systems, August, 1995
October 1995 (11 pages)

The Bay Area Gigabit Network Testbed (BAGNet) is a high-performance ATM (155 Mbps) testbed located within the San Francisco Bay Area in northern California. BAGNet is a metropolitan-area network, spanning an area of approximately 50 square miles. There are fifteen sites participating in the testbed, with up to four hosts per site. Although BAGNet is an applications-oriented testbed, much of our effort has been directed towards getting the testbed running and understanding the factors that impact performance of an ATM network. We present some of our experiences in this paper.

TR 95.22

AN EDGE-BASED SOLUTION-ADAPTIVE METHOD APPLIED TO THE AIRPLANE CODE

Rupak Biswas, Scott D. Thomas (Sterling Software Inc.) and Susan E. Cliff (NASA Ames)
November 1995 (11 pages)

Computational methods to solve large-scale realistic problems in fluid flow can be made more efficient and cost effective by using them in conjunction with dynamic mesh adaption procedures that perform simultaneous coarsening and refinement to capture flow features of interest. This work couples the tetrahedral mesh adaption scheme, called 3D_TAG, with the AIRPLANE code to solve complete aircraft configuration problems in transonic and supersonic flow regimes. Results indicate that the near-field sonic boom pressure signature of a cone-cylinder is improved, the oblique and normal shocks are better resolved on a transonic wing, and the bow shock ahead of an unstarted inlet is better defined.

TR 95.23

A DETERMINISTIC PARTICLE METHOD FOR ONE-DIMENSIONAL REACTION-DIFFUSION EQUATIONS

Michael Mascagni (Center for Computing Sciences, I.D.A.)
November 1995 (15 pages)

We derive a deterministic particle method for the solution of nonlinear reaction-diffusion equations in one spatial dimension. This deterministic method is an analog of a Monte Carlo method for the solution of these problems that has been previously investigated by the author. The deterministic method leads to the consideration of a system of ordinary differential equations for the positions of suitably defined particles. We then consider the time explicit and implicit methods for this system of ordinary differential equations and we study a Picard and Newton iteration for the solution of the implicit system. Next we solve numerically this system and study the discretization error both analytically and numerically. Numerical computation shows that this deterministic method is automatically adaptive to large gradients in the solution.

TR 95.24

AERODYNAMIC DESIGN OPTIMIZATION VIA REDUCED HESSIAN SQP WITH SOLUTION REFINING

Dan Feng and Thomas H. Pulliam (NASA Ames Research Center)
17 pages, Dec 1995.

An all-at-once reduced Hessian Successive Quadratic Programming (SQP) scheme has been shown to be efficient for solving aerodynamic design optimization problems with a moderate number of design variables. This paper extends this scheme to allow solution refining. In particular, we introduce a reduced Hessian refining technique that is critical for making a smooth transition of the Hessian information from coarse grids to fine grids. Test results on a nozzle design using quasi-one-dimensional Euler equations show that through solution refining the efficiency and the robustness of the all-at-once reduced Hessian SQP scheme are significantly improved.

TR 95.25

THE EFFECT OF NEURAL NETWORK STRUCTURE ON A MULTISPECTRAL LAND-USE CLASSIFICATION

Justin D.Paola and Robert A. Schowengerdt (University of Arizona, Tucson)
December 1995 (31 pages)

While neural networks are now an accepted alternative to statistical multispectralclassification techniques for remote sensing image classification, the network approach presents both unique challenges and abilities. The size of the hidden layer must be determined by trial and error, and the random initial weight settings result in different paths for the training procedure, making the network a non-deterministic classifier. For the sample classification presented here, it was found that there was a range of optimal hidden layer sizes below which the accuracy decreased and above which the training time increased. However, it was also found that for a fairly wide range, the hidden layer size made little difference to the final classification accuracy. Initial weight randomization was as much of a factor as hidden layer size. Using 3x3 windows of data in each band was found, despite increased training time per iteration, to achieve similar accuracy with less overall training time, although with less consistency.

IV. PUBLICATIONS

RUPAK BISWAS

Strawn, R. C. (US Army AFDD), Biswas, R., Lyrintzis, A. S. (Purdue), "Helicopter Noise Predictions using Kirchhoff Methods," Proceedings of the 51st AHS Annual Forum, Fort Worth, TX, pages 495-508, May 1995.

Biswas, R., and Strawn, R. C. (US Army AFDD), "Anisotropic h-Refinement for Unstructured Hexahedral Grids," Proceedings of the 3rd U.S. National Congress on Computational Mechanics, Dallas, TX, page 234, June 1995.

Duque, E. P. N. (US Army AFDD), Biswas, R., and Strawn, R. C. (US Army AFDD), "A Solution Adaptive Structured/Unstructured Overset Grid Flow Solver with Applications to Helicopter Rotor Flows," 13th AIAA Applied Aerodynamics Conference, San Diego, CA, AIAA Paper 95-1766, June 1995.

Biswas, R., and Dagum, L. (Silicon Graphics), "Parallel Implementation of an Adaptive Scheme for 3D Unstructured Grids on a Shared-Memory Multiprocessor," Proceedings of Parallel CFD'95 Conference, Pasadena, CA, June 1995.

Strawn, R. C. (US Army AFDD), and Biswas, R., "Computation of Helicopter Rotor Acoustics in Forward Flight," Journal of the American Helicopter Society Volume 40, pages 66-72, July 1995.

Strawn, R. C. (US Army AFDD), Biswas, R., and Garceau, M. (Stanford), "Unstructured Adaptive Mesh Computations of Rotorcraft High-Speed Impulsive Noise," Journal of Aircraft Volume 32, pages 754-760, July/August 1995.

Biswas, R., and Strawn, R. C. (US Army AFDD), "Dynamic Mesh Adaption for Tetrahedral Grids," Lecture Notes in Physics, Volume 453, pages 127-132, October 1995.

Strawn, R. C. (US Army AFDD), and Biswas, R., "Numerical Simulations of Helicopter Aerodynamics and Acoustics," to appear in Journal of Computational and Applied Mathematics.

Biswas, R., and Strawn, R. C. (US Army AFDD), "Mesh Quality Control for Multiply-Refined Tetrahedral Grids," to appear in Applied Numerical Mathematics.

Biswas, R., and Strawn, R. C. (US Army AFDD), "A Dynamic Mesh Adaption Procedure for Unstructured Hexahedral Grids," accepted to 34th AIAA Aerospace Sciences Meeting, Reno, NV.

Duque, E. P. N. (US Army AFDD), Strawn, R. C. (US Army AFDD), Biswas, R., and Ahmad, J. U. (Sterling), "An Overset Grid Navier-Stokes/Kirchhoff-Surface Method for Rotorcraft Aeroacoustic Predictions," accepted to 34th AIAA Aerospace Sciences Meeting, Reno, NV.

RUPAK BISWAS (cont'd)

Wissink, A. M. (Univ. of Minnesota), Lyrintzis, A. S. (Purdue), Biswas, R., Oliker, L., and Strawn, R. C. (US Army AFDD), "Efficient Helicopter Aerodynamic and Aeroacoustic Predictions on Parallel Computers," accepted to 34th AIAA Aerospace Sciences Meeting, Reno, NV.

Biswas, R., Thomas, S. D. (Sterling), and Cliff, S. E., "An Edge-Based Solution-Adaptive Method Applied to the AIRPLANE Code," accepted to 34th AIAA Aerospace Sciences Meeting, Reno, NV.

Strawn, R. C. (US Army AFDD), Oliker, L., DeRyke, D. (Sterling), and Biswas, R., "New Computational Methods for the Prediction and Analysis of Helicopter Acoustics," submitted to 2nd AIAA/CEAS Aeroacoustics Conference, State College, PA.

Sohn, A. (NJIT), Biswas, R., and Simon, H. D. (Silicon Graphics), "Impact of Load Balancing on Unstructured Adaptive Grid Computations for Distributed-Memory Multiprocessors," submitted to 10th ACM International Conference on Supercomputing, Philadelphia, PA.

Biswas, R., and Strawn, R. C. (US Army AFDD), "Dynamic Mesh Adaption for Unstructured Hexahedral Grids," submitted to 15th International Conference on Numerical Methods in Fluid Dynamics, Monterey, CA.

Sohn, A. (NJIT), Biswas, R., and Simon, H. D. (Silicon Graphics), "Dynamic Load Balancing for Unstructured Adaptive Grids," submitted to SIAM Annual Meeting, Kansas City, MO.

DAN FENG

D. Feng and T. H. Pulliam, "An All-at-Once Reduced Hessian SQP Scheme for Aerodynamics Design Optimization, TR 95.19," Presented at SIAM Annual meeting, October 1995.

D. Feng and T. H. Pulliam, "Aerodynamic Design Optimization via Reduced Hessian SQP with Solution Refining, TR 95.24," Submitted to 27th AIAA Fluid Dynamics Conference.

D. Feng, R. B. Schnabel, "Local Convergence Analysis of Tensor and SQP Methods for Singular Constrained Optimization, TR 95.10," Submitted to SIAM Journal on Optimization.

R. H. Byrd, D. Feng and R. B. Schnabel, "On Optimality Conditions for Singular Constrained Optimization, TR 95.03," Submitted to Mathematical Programming.

D. Feng and R. B. Schnabel, "Tensor Methods for Equality Constrained Optimization," To appear in SIAM Journal on Optimization.

MARJORY J. JOHNSON

Marjory J. Johnson, ``Experiences with the Bay Area Gigabit Network Testbed," Proceedings of the 5th IEEE Workshop on Future Trends in Distributed Computer Systems, Aug. 1995, pp. 26-32 (also RIACS TR #95.21).

Marjory J. Johnson and Jeffrey N. Townsend, ``Achieving High Throughput for Data Transfer over ATM Networks," accepted for presentation at the International Communications Conference, ICC '96.

JOSEPH OLIGER

Gustafsson, B., Kreiss, H.-O., Oliger, J., "Time Dependent Problems and Difference Methods, to be published by J. Wiley and Sons in 1996.

PELLE OLSSON

Gustafsson, B. (Uppsala Univ.), Olsson, P., "Fourth-Order Difference Methods for Hyperbolic IBVP," Journal of Computational Physics, volume 117, March 1995, pages 300-317.

Olsson, P., "Summation by Parts, Projections, and Stability I, Mathematics of Computations," volume 64, number 211, July 1995, pages 1035-1065.

Olsson, P., "Summation by Parts, Projections, and Stability II, Mathematics of Computations," volume 64, number 212, October 1995, pages 1473-1493.

BARNEY PELL

B. Pell., "A Strategic Metagame Player for General Chess-Like Games, Computational Intelligence," 12(1), 1996 (To appear).

Barney Pell, Susan L. Epstein (Hunter College, SUNY) and Robert Levinson (UCSC), "Introduction to the Special Issue on Games: Structure and Learning. Computational Intelligence," 12(1), 1996 (To appear).

C. Baudin (Recom), S. Kedar (ILS), and B. Pell, "Increasing Levels of Assistance in Refinement of Knowledge-Based Retrieval Systems," In Gheorghe Tecuci and Yves Kodratoff (eds), "Machine Learning and Knowledge Acquisition: Integrated Approaches", Academic Press, 1995.

Muscettola N. (Recom), Pell, B., Hansson, O. (HRI), and Mohan, S. (RECOM), In Henry, G.W. and Eaton J.A. (eds), Automating Mission Scheduling for Space-Based Observatories, "Robotic Telescopes: Current Capabilities, Present Developments, and Future Prospects for Automated Astronomy," ASP Conf. Ser. No. 79, Astronomical Society of the Pacific, Provo, UT, 1995.

BARNEY PELL (cont'd)

Barney P., Bernard, D. E., Chien, S. A., Gat, E., Muscettola, N., P. Pandurang Nayak, Wagner, M.D., and Williams, B. C., "An implemented architecture integrating onboard planning, scheduling, execution, diagnosis, monitoring and control for autonomous spacecraft," submitted to the National Conference on Artificial Intelligence, 1996.

JAMES REUTHER/ANTONY JAMESON

Reuther, J. and Jameson, A., "Aerodynamic Shape Optimization of Wing and Wing-Body Configurations Using Control Theory," 33rd Aerospace Sciences Meeting and Exhibit, Reno, Nevada, AIAA paper, 95-0123, January 1995.

Reuther, J. and Jameson, A., "A Comparison of Design Variables for Control Theory Based Airfoil Optimization," In proceedings of the 6th International Symposium on Computational Fluid Dynamics, Lake Tahoe, Nevada, September 1995.

Reuther, J. and Jameson, A., "Supersonic Wing and Wing-Body Shape Optimization Using an Adjoint formulation," In proceedings of The Forum on {CFD} for Design and Optimization, (IMECE 95), San Francisco, California, November 1995.

Gallman, J. and Reuther, J. and Pfeiffer, N. and Forrest, W. and Bernstorf, D., "Business Jet Wing Design Using Aerodynamic Shape Optimization," 34th Aerospace Sciences Meeting and Exhibit, to appear, Reno, Nevada, AIAA paper, 96-0554, January 1996.

Reuther, J. and Jameson, A. and Farmer, J. and Martinelli, L., "Aerodynamic Shape Optimization of Complex Aircraft Configurations via an Adjoint Formulation," 34th Aerospace Sciences Meeting and Exhibit, Reno, Nevada, AIAA paper, 96-0094, to appear, January 1996.

ROBERT SCHREIBER

Chatterjee, S., Gilbert, J. R. (Xerox), Schreiber, R, and Sheffler, T. J., "Modeling Data-Parallel Programs with the Alignment-Distribution Graph," Journal of Programming Languages, Special issue on compiling and run-time issues for distributed address space machines, to appear.

Chatterjee, S., Gilbert, J. R. (Xerox PARC), Long, F. J. E. (UCSC), Schreiber, R., and Teng, S.-H. (Xerox PARC), "Generating Local Addresses and Communication Sets for Data-Parallel Programs," Journal of Parallel and Distributed Computing, JPDC, Vol 26, Num 1, pages 72-84, April 1995.

Chatterjee, S., Gilbert, J. R. (Xerox PARC), Schreiber, R., and Teng, S.-H. (Xerox PARC), "Optimal Evaluation of Array Expressions on Massively Parallel Machines," ACM Trans actions on Programming Languages and Systems, Toplas, Vol. 17, Num 1, pages 123-156, January 1995.

ROBERT SCHREIBER (cont'd)

Chatterjee, S., Gilbert, J. R. (Xerox PARC), Schreiber, R., and Sheffler, T. J., "Array Distribution in Data-Parallel Programs," Languages and Compilers for Parallel Computing, Lecture Notes in Computer Science Series. In K. Pingali, U. Banerjee, D. Gelernter, A. Nicolau, and D. Padua, editors, Languages and Compilers for Parallel Computing, LNCS Vol 892, Springer-Verlag, 1995, pp. 76--91.

Sheffler, T. J., Schreiber, R., Gilbert, J. R. (Xerox PARC), and Chatterjee, S., "Aligning Parallel Arrays to Reduce Communication," Proceedings of Frontiers '95, McLean, VA, to appear February 1995.

Bailey, D. H., Bjorstad, P. E., Gilbert, J. R., Mascagni, M.V., Schreiber, R.S., Simon, H.D., Torczon, V. J., and Watson, L.T., co-editors, "Proceedings of the Seventh SIAM Conference on Parallel Processing for Scientific Computing." SIAM, 1995.

Chong, F., and Schreiber, R., "Parallel Sparse Triangular Solution with Partitioned Inverse and Prescheduled DAGs," Workshop on Solving Irregular Problems on Distributed Memory Machines, 1995.

Rothberg, E. and Schreiber, R., "Efficient Parallel Sparse Cholesky Factorization," Proceedings of the Seventh SIAM Conference on Parallel Processing for Scientific Computing, SIAM, 1995.

Schreiber, R., with Chong, F., "Parallel Sparse Triangular Solution with Partitioned Inverse and Prescheduled DAGs." Proceedings Workshop on Solving Irregular Problems on Distributed Memory Machines, 1995.

Schreiber, R., with Sheffler, T, J., Pugh, W., Gilbert, J. R., and Chatterjee, S., "Efficient Distribution Analysis Via Graph Contraction," submitted to the Eight Workshop on Languages and Compilers for Parallel Computing.

Schreiber, R., with Sheffler, T, J., Gilbert, J. R., and Chatterjee, S., "Algorithms For Automatic Alignment of Arrays," submitted to Parallel and Distributed Computing.

THOMAS J. SHEFFLER

Sheffler, T. J., Schreiber, R., Gilbert, J. R. (Xerox PARC), and Chatterjee, S., "Aligning Parallel Arrays to Reduce Communication," Proceedings of Frontiers '95, McLean, VA, to appear February 1995.

Sheffler, T. J. and Chatterjee, S., "An Object-Oriented Approach to Nested Data Parallelism," Proceedings of Frontiers '95, the Fifth Symposium on the Frontiers of Massively Parallel Computation, McLean, VA, February 1995.

WEI-PAI TANG

Tang, W.-P. (Univ. of Waterloo), S. Clift, (Queens University), "Weighted Graph Based Ordering Techniques for Preconditioned Conjugate Gradient Methods" {*BIT*}. Vol. 35, No. 1, pp. 30–47, 1995.

Tang, W.-P. (Univ. of Waterloo), George, J. A. (Univ. of Waterloo), Ikramov, K. and A. N. Krivoshapova, (Univ. of Moscow), SIAM J. Matrix Analysis and Applications, Vol. 16, No. 2, pp. 377-387, 1995.

Tang, W.-P. (Univ. of Waterloo), George, J. A. (Univ. of Waterloo), Ikramov, K. and L. Matushkina, (Univ. of Moscow), "On a QR-like Algorithm for Some Structured Eigen value Problems," SIAM J. Matrix Analysis and Applications, Vol. 16, No. 4, pp 1107-1126 1995.

Fan, Q., Forsyth,P. A., Univ. of Waterloo, McMacken, J. R. F. (Goal Electronics Inc.), Tang, W.-P. (Univ. of Waterloo), "Performance Issues for Iterative Solvers in Device Simulation," accepted by SIAM J. Scientific Computing.

Tang, W.-P. (Univ. of Waterloo), and H. Sun (Univ. of Waterloo), "Overdetermined Schwarz alternating method," accepted by SIAM Journal on Scientific Computing,

Tang, W.-P. (Univ. of Waterloo), J. Oliger, L. C. Huang and Y. D. Wu, "Toward effective and robust finite difference schemes for unsteady incompressible Navier-Stokes Equation on the half-staggered mesh," Sixth International Symposium of CFD, Lake Tahoe, CA. U.S.A. September 1995, 5 pages.

Tang, W.-P. (Univ. of Waterloo), and H. Sun "An overdetermined Schwarz alternating method," International Symposium on Domain Decomposition Methods for Partial Differential Equations, SIAM, 1995,

Tang, W.-P. (Univ. of Waterloo), S. Clift (Queens University) and H. Simon (SGI), "Spectral ordering techniques for ILU preconditioners," submitted to IMA Journal of Numerical Analysis. 24 pages

Tang, W.-P. (Univ. of Waterloo), G. Golub (Stanford Univ), L. C. Huang (ICMSEC) and H. Simon (SGI) "A fast solver for incompressible Navier-Stokes equations with finite difference methods," submitted to SIAM Journal on Scientific Computing, 17 pages.

V. SEMINARS AND COLLOQUIA

RUPAK BISWAS

Anisotropic h-Refinement for Unstructured Hexahedral Grids, Biswas, R., 3rd U.S. National Congress on Computational Mechanics, Dallas, TX, June 13, 1995.

Parallel Implementation of an Adaption Scheme for 3D Unstructured Grids on a Shared Memory Multiprocessor, Biswas, R., Parallel CFD'95 Conference, Pasadena, CA, June 26, 1995.

New Computational Models for Helicopter Aerodynamics and Acoustics, Biswas, R., Brown University, Providence, RI, October 20, 1995.

Dynamic Mesh Adaption for Unstructured Hexahedral Grids, Biswas, R., SIAM Annual Meeting, Charlotte, NC, October 24, 1995.

FREDRIC CHONG

Building a Better Butterfly: The Multiplexed Metabutterfly, Chong, F., NASA Ames Research Center, March 24, 1995.

DAN FENG

Tensor-Krylov Methods for Large Nonlinear Equations, Analogy, Inc., Beaverton, Oregon, November 2, 1995.

An All-at-Once Reduced Hessian SQP Scheme for Aerodynamics Design Optimization, SIAM Annual Meeting, Charlotte, North Carolina, October 24, 1995.

Tensor-Krylov Methods for Large Nonlinear Equations, Scientific Computation and Computational Mathematics Seminar, Stanford University, May 15, 1995.

An All-at-Once Reduced Hessian SQP Scheme for Aerodynamics Design Optimization, Boeing Computer Services, Seattle, Washington, March 15, 1995.

GARY L. MILLER

A Delaunay Based Numerical Method for Three Dimensions: Generation, Formulation, and Partition, Miller, G.L., NASA Ames Research Center, August 24, 1995.

ANDREW SOHN

Dynamic Load Balancing for Adaptive Unsteady CFD on an SP-2 Multiprocessor, Sohn, A., NASA Ames Research Center, August 3, 1995.

WEI-PAI TANG

Invited lecture: "A Hybrid Parallel Preconditioning Algorithm for CFD," Parallel Computational Fluid Dynamics 95, California Institute of Technology, Pasadena, CA. June 26-29, 1995.

Invited lecture: "Toward Efficient and Robust Finite Difference Schemes for Unsteady Incompressible Navier-Stokes Equation - on the Half Staggered Mesh," Sixth International Symposium on CFD, Lake Tahoe, CA. U.S.A., September 1995.

NASA Ames Research Center, NAS Seminar. "High Performance Preconditioners," January 1995.

Stanford Univ., Scientific Computing and Computational Mathematics Seminar, "Effective Sparse Approximate Inverse Preconditioners," January 1995.

Dept. of Mathematics, UCLA, "Effective Sparse Approximate Inverse Preconditioners," February 1995.

Dept. of Applied Mathematics, Univ. of Washington, "A Hybrid Parallel Preconditioner," November 1995.

Boeing Information and Support Services, "Hybrid Parallel Preconditioners in CFD," November 1995.

MARCO ZAGHA

Efficient Irregular Computation on Pipelined-Memory Multiprocessors, Zagha, M., NASA Ames Research Center, June 6, 1995.

VI. OTHER ACTIVITIES

RUPAK BISWAS

Attended and presented a paper at the 3rd U.S. National Congress on Computational Mechanics, Dallas, TX, June 12-14, 1995.

Attended and presented a paper at the Parallel CFD'95 Conference, Pasadena, CA, June 26-28, 1995.

Visited Professor Chi-Wang Shu and gave a talk at Brown University, Providence, RI, October 19-22, 1995.

Attended and presented a paper at the SIAM Annual Meeting, Charlotte, NC, October 23-26, 1995.

Hosted Professor Andrew Sohn from New Jersey Institute of Technology, Newark, NJ, July 4 - August 5, and December 26-31, 1995.

MARJORY J. JOHNSON

Chaired session at 5th IEEE Workshop on Future Trends in Distributed Computer Systems, Korea, August 1995.

Presented invited paper at 5th IEEE Workshop on Future Trends in Distributed Computer Systems, Korea, August 1995.

Co-coordinator for BAGNet testbed.

Participated in several gigabit testbed workshops.

Member, ISO/TC20/USSCAG13 committee to develop communication standards for space missions.

Member, program committee for 6th IFIP International Conference on High Performance Networking '95, Spain, Sept. 1995.

Member, program committee for COST 237 Multimedia Conference, Copenhagen, Denmark, Nov. 1995.

JOSEPH OLIGER

Professor of Computer Science, Stanford University.

Member, Board of Trustees of the Society for Industrial and Applied Mathematics.

Associate Editor, Journal of Computational and Applied Mathematics.

Editorial Advisory Committee, SIAM Review.

BARNEY PELL

Edited Computational Intelligence Journal, Special Issue on Games: Structure and Learning 12(1), 1996 (To appear).

Reviewer for Journal of Artificial Intelligence Research (JAIR).

AAAI Spring Symposium on Lessons Learned from Implemented Software Architectures for Physical Agents, Stanford, March 1995.

International Joint Conference on Artificial Intelligence, Montreal, Quebec, August 1995 (IJCAI-95).

Workshop on Agent Theories, Languages, and Architectures, at IJCAI-95.

AWARD:
Contractor Certificate of Excellence (Team Award), Ames Contractor Council, October 1995. For development of New Millennium Autonomy Architecture Prototype.

JAMES REUTHER

Worked with Code AAH in the High Speed Research Project to apply optimization techniques in the practical development of the next generation High Speed Civil Transport.

Participated in several wind tunnel test to support the HSR project that validated optimized designs.

Participated in dozens of teleconferences with the Configuration Aerodynamics team for the HSR project to discuss various aspects of optimized configurations and determine future technology development.

ROBERT S. SCHREIBER

Minisymposium organizer, Intl. Conf. on Industrial and Applied Mathematics, Hamburg, Germany, July 1995.

Organizer, User-interface Workshop on Parallel Programming, Santa Barbara, CA, 1995

General Chair, ACM Principle and Practices of Parallel Programming Conference, to be held June 1997.

Guest editor (with Piyush Mehrotra), special issue of "Scientific Programming" on implementation of HPF.

SIGNUM Vice-chair.

Co-instructor (with A. Edelman and S-h. Teng) in course 6.50s "A Peek at Parallel Processing from an Applications Perspective", MIT Summer Professional Program, June 1995.

Ph.D. thesis committee of Susan Hinrichs, Carnegie-Mellon Univ.

Attended SIG Officers Orientation Meeting, ACM HQ, NY, October, 1995.

Research collaboration with Prof. R. Street on Stanford NSF grand challenge project.

STEVEN SUHR

Contributed a new scrolling algorithm which has been accepted for incorporation into a future version of the text editor GNU emacs, improving the visual stability of the approach which emacs uses in updating the displayed text on the screen.

VII. RIACS STAFF

ADMINISTRATIVE STAFF

Joseph Oliger, Director - Ph.D., Computer Science, University of Uppsala, Sweden, 1973. Numerical Methods for Partial Differential Equations (03/25/91 - present).

Frances B. Abel, Office and Financial Manager (5/5/88 - present).

Deanna M. Gearhart, Administrative Assistant II (5/9/88 - present).

Rufus White Jr., Systems Administrator (5/17/93 - present).

RIACS SCIENCE COUNCIL

Dr. Robert B. Schnabel (Convener), Department of Computer Science, University of Colorado, Boulder, Colorado 80309-0430, bobby@cs.colorado.edu.

Dr. James W. Demmel, Computer Science Division, 571 Evans Hall, University of California, Berkeley, California 94720, demmel@cs.berkeley.edu.

Dr. David Gottlieb, Division of Applied Mathematics, Brown University, Providence, Rhode Island 02912, am420000@brownvm.brown.edu.

Dr. Joseph Monroe, North Carolina A&T University, Greensboro, North Carolina, 27411, monroe@garfield.nest.edu.

Dr. Kenneth W. Neves, Boeing Company, P.O. Box 24346, MS 7L-25, Seattle, Washington 98124, kneves@atc.boeing.com.

Dr. Joseph Oliger (Ex-Officio), Dir., Research Institute for Advanced Computer Science, NASA Ames Research Center, Moffett Field, California 94035-1000, oliger@riacs.edu.

Dr. Thomas H. Pulliam, NASA Ames Research Center, MS 202A-1, Moffett Field, CA 94035, pulliam@rft29.nas.nasa.gov.

Dr. Daniel A. Reed, Department of Computer Science, 1304 West Springfield Avenue, University of Illinois, Urbana, Illinois 61801-2987, reed@cs.uiuc.edu.

Dr. Marc Snir, IBM, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598, snir@watson.ibm.com.

Dr. David Cummings, Executive Director, Universities Space Research Association, Columbia, Maryland 21044, cummings@hq.usra.edu.

SCIENTIFIC STAFF

Rupak Biswas, Ph.D., Computer Science, Rensselaer Polytechnic Institute, 1991, Large scale scientific computation using parallel and adaptive methods (9/16/91 - present).

Wray Buntine, Ph.D., Computer Science, University of Technology, Sydney, Australia, 1992, Mathematical and probabilistic modeling of problems in intelligent systems (2/5/90 - 5/31/95).

Peter Cheeseman, Ph.D., Artificial Intelligence, Monash University, Australia 1979, Artificial intelligence and automatic control, induction of models under uncertainty, Bayesian inference, expert systems and robotics (5/6/85 - 6/30/95).

Dave Gehrt, JD Law, University of Washington, 1972, UNIX system administration, security, and network based tools (1/84 - 7/85, 2/1/88 - present).

Marjory J. Johnson, Ph.D., Mathematics, University of Iowa, 1970, High-performance net working for both space and ground applications (1/9/84 -present).

Robert S. Schreiber, Ph.D. - Computer Science, Yale University, 1977, Parallel numerical algorithms and parallel computer architectures (8/29/88 - 12/8/95).

VISITING SCIENTISTS

Marsha Berger, Ph.D. - New York University, Computational fluid dynamics; parallel computing (6/26/95 - 8/24/95).

Alfred Brennis, Ph.D. - Researcher, Daimler-Benz, Germany, Dynamics of numerics for reacting flows (5/3/95 - 9/30/95).

Tony F. Chan, Ph.D. - Professor of Mathematics, University of California, Los Angeles, Efficient algorithms in large-scale scientific computing, parallel algorithms and computational fluid dynamics (8/14/95 - 8/30/95).

Herman Deconick, Ph.D. - Professor, von Karmen Institute, Belgium, CFD Algorithm development for compressible flow.

Bertil Gustafsson, Ph.D. - Professor, Upsalla University, Sweden, Numerical methods for partial differential equations. Computational fluid dynamics (2/18/95 - 3/3/95).

Antony Jameson, Ph.D. - McDonnell Professor of Aerospace Engineering, Princeton University, Numerical Methods, computational fluid dynamics, computational sciences (2/1/95 - 8/31/95).

Michael Masgagni, Ph.D. - Researcher, Center for Computing Sciences, Parallel computing, numerical analysis, Monte Carlo Methods (6/27/95 - 9/1/95).

Robert Schowengerdt, Ph.D. - Professor, University of Arizona, Image processing for earth remote sensing/Pattern recognition in large image databases/remote sensor modeling and simulation (6/5/95 - 7/20/95).

VISITING SCIENTISTS (cont'd)

Andrew Sohn, Ph.D. - Assistant Professor, New Jersey Institute of Technology, Dynamic load balancing for grid partitioning on SP-2 (7/5/95 - 8/4/95).

Peter Sweby, Ph.D. - Lecturer in Mathematics, University of Reading, England, The dynamics of discretisations of convection reaction diffusion equations was investigated over a range of physical and numerical parameters (7/31/95 - 8/18/95).

Stephen Vavasis, Ph.D. - Associate Professor, Cornell University, Finite-element mesh generation sparse matrix computation numerical optimization (6/1/95 - 8/20/95).

David Zingg, Ph.D. - Associate Professor, University of Toronto, Canada, Development and analysis of high-accuracy numerical methods applicable to simulations of fluid Flows, Acoustic Waves, and Electromagnetic Waves (6/27/94 - 9/2/94).

POST-DOCTORAL SCIENTISTS

Dan Feng, Ph.D. - Computer Science, University of Colorado, Boulder, Numerical computation including optimization, solving systems of nonlinear equations, solving PDE's and their applications in aeronautics and astronautics; parallel numerical computation (9/1/93 - 12/8/95).

Pelle Olsson, Ph.D. - Uppsala University, Sweden, Initial-boundary value problems for hyperbolic and parabolic PDEs, numerical methods for PDEs on parallel computers (11/2/92 - 1/3/95).

Barney Pell, Ph.D. - University of Cambridge, England, Artificial intelligence, machine learning, strategic reasoning, logic programming, automatic scheduling, information retrieval, multiple autonomous agents, artificial life (9/27/93 - 6/30/95).

Thomas Scheffler, Ph.D. - Carnegie Mellon, Computer engineering, Parallel compiler techniques (6/11/93 - 7/10/95).

RESEARCH ASSOCIATES

Leonid Oliker - Computer Science, University of Colorado, compilation of data parallel programs (9/1/94 - present).

James Reuther, Ph.D. - University of California, Davis, numerical optimization aerodynamic shape optimization numerical analysis CFD(9/6/94 - present).

Steven Suhr, - Computer Science, Stanford University, programming languages (7/1/92 - present).

STUDENTS

Ender E. Bilir - Computer Science, Princeton University, Multiprocessor compilers, algorithms, (6/5/95 - 9/5/95)

James Dougherty - Computer Science, Stanford University, Parallel and distributed computing, networks, machine learning, (6/21/95 -8/4/95)

Susie Go - Applied Mathematics, University of California, Los Angeles, Grid generation and coarsening techniques for unstructured methods, (6/28/95 - 9/8/95).

Justin Paola - Electrical & Computer Engineering, University of Arizona, image processing/remote sensing, University of Arizona (5/30/92 - 9/7/95).

Jeffrey Townsend - Computer Science, Stanford University, Research for Bay Area Gigabit ATM network testbed, 6/21/95 - present)

Xiaolei Zhu - Computer Science, Stanford University, error analysis for adaptive grids (6/19/95 - 9/20/95).

CONSULTANTS

Saul Abarbanel - Professor, Tel-Aviv University, Numerical methods for partial differential equation (8/18/95 - 8/19/95).

Marsha Berger, Ph.D. - New York University, Computational fluid dynamics; parallel computing (1/1/93 - present).

Tony F. Chan - Professor of Mathematics, University of California, Los Angeles, Efficient algorithms in large-scale scientific computing, parallel algorithms and computational fluid dynamics (10/01/86 - present).

Robert Coreless - Professor, University of Western Ontario, Canada, Numerical simulations of nonlinear differential equations (1/24/95 - 1/26/95).

Christopher Diot - Research Scientist, INRIA, High-speed networking, multimedia (8/21/95 - 8/22/95).

John Gilbert - Research Scientist, Xerox Palo Alto Research Center, Parallel computing and theoretical computer science (5/1/92 -12/31/95).

Richard G. Johnson, Ph.D. - Physics, Indiana University, 1956, Global environmental problems and issues (11/1/92 - present).

Gary Miller - Professor, Carnegie Mellon University, Computational geometry, parallel scientific computation (8/21/95 - 8/24/95).

CONSULTANTS (cont'd)

Pelle Olsson, Ph.D. - Assistant Professor, Stanford University, Initial-boundary value problems for hyperbolic and parabolic PDEs, numerical methods for PDEs on parallel computers (1/4/95 - present).

William Pugh - Associate Professor, University of Maryland, Excalibur project on techniques for automatic placement of data and computations on parallel computers (10/1/94 - 3/31/95).

Robert Schnabel, Ph.D. - Professor, University of Colorado, Boulder, Numerical computation especially optimization nonlinear equations, parallel computation (1/1/94 - present).

Horst Simon - Researcher, Silicon Graphics, Inc. (7/1/95 - present).

Wei-Pai Tang, Ph.D. - Professor, University of Waterloo, Canada, Numerical solution of partial differential equations, numerical linear algebra, parallel computations (7/1/94 - present).

RIACS
Mail Stop T041-5
NASA Ames Research Center
Moffett Field, CA 94035